

2018-01-20

Tom Rochette <tom.rochette@coreteks.org>

November 2, 2024 — [36c8eb68](#)

0.1 Context

0.2 Learned in this study

0.3 Things to explore

1 Problems faced

2 Overview

2.1 kNN

- Given a set of data points where some of the point's coordinates are inputs and some are outputs, is it possible for us to predict, given a data point with missing coordinates, those missing values?
- Let's start with a simple case, where we have the following points: $(0, 0)$, $(0, 1)$, $(1, 0)$, $(1, 1)$
- If we are given the following point $(0.5, ?)$, how could we determine the missing value?
 - Let say we group together (cluster) all the points which have the same x values, such that $(0, 0)$ and $(0, 1)$ are together and $(1, 0)$ and $(1, 1)$ are together. If we average these clusters, we get two points, $(0, 0.5)$ and $(1, 0.5)$. In this case, none of the two points can help us decide what is the missing value.
 - If we decide to cluster all points which have the same y values, we get a different result: $(0.5, 0)$ and $(0.5, 1)$. Here we have a dilemma, as both values have the same x, thus it is not possible to decide which one to pick
 - We can compute the center point of the four points of our initial case, which gives us the centroid $(0.5, 0.5)$. Here, the answer would be 0.5. But let say that our y value is supposed to be discrete (because it represents a category), what should the value be?
 - * In this case, since we have 2 points with a value of 0 and 2 points with a value of 1, there is 50% probability that it belongs to either of these values. If there was a third point, for instance $(0.5, 0)$, then this would tip the balance in favor of 0 being the suspected value.

2.2 Reinforcement learning

- We have an agent within an environment. This agent can make observations and then take actions. The goal of this agent is to accomplish goals. If we are to give many goals to this agent, then we need a way to communicate whether a certain goal has been completed or not. In the reinforcement learning literature, this is known as the reward signal.
- When an agent takes an action, this action can have a few different results: bringing the agent closer to its goal (+1), keeping the agent the same distance from its goal (0), or moving the agent away from its goal (-1). While an agent is trying to reach its goal, it may receive a signal that it is making progress toward its goal, but it may also receive no such indication. In this latter case, it is much more difficult for the agent to make progress because there is nothing to follow. In the former case, the agent will want to follow the path that provides positive rewards.

- Sometimes the idea of being at the same distance of its goal can be seen as a bad thing. In simple examples, called grid world, we have an agent that can travel through a grid where there is a reward signal. In its simplest case, the agent starts on a square which provides reward signal = 0, and on its right, a -1 signal. In this instance, it is better for the agent not to move to the next square. In an extended version of this problem

3 See also

4 References

- https://en.wikipedia.org/wiki/K-nearest_neighbors_algorithm